

## WE ALSO NEED COMPLETE BEHAVIORAL MODELS

WILLIAM R. HUTCHISON

BEHAVIOR SYSTEMS, LLC

I am very pleased to see Donahoe, Palmer, and Burgos foster this important discussion with a very clear and cogent paper. I find their reasoning inescapable that environment-behavior relations are what is always conditioned. My own computer model of operant behavior is built on the same assumptions (Hutchison, 1984, 1985, 1995, 1997; Stephens & Hutchison, 1993), and the many simulations I have done are uniformly supportive of their claim that a single model can produce discriminated and undiscriminated conditioning in a wide variety of situations (e.g., Hutchison, 1997). I hope I can now come out of the closet without having to wear a scarlet S-R on my chest.

I also strongly agree with Donahoe et al. regarding the value of computer modeling as a scientific tool. Donahoe and Palmer (1994) describe computer simulations as formal interpretation (p. 128), automatically "reasoning" or predicting from a set of assumptions to their implications. All scientists are constantly making predictions from theoretical formulations, both in designing experiments and in judging the congruence of data with predictions from various formulations. Just as we use a computer spreadsheet to ask "what if?" questions when data get too complicated to manage on paper, we can use a computer model to ask "what if?" questions of complicated behavioral relations. Our subject matter concerns environment-behavior systems that complicate enormously as the history of the organism develops. In my opinion, computer models are essential for further progress in our discipline, and they imply no necessary compromise of our radical behavioral philosophy of science. Quantitative behavior analysts are increasingly using these tools in valuable ways; for example, simulating how molecular conditioning processes that occur

at each reinforcing event in variable schedules can account for some of the deviations from the matching law at a molar level. The analyses could in principle have been done on paper, but no scientist would take the time to do so.

A common misunderstanding is that computer modelers are substituting computer data for real animal data. On the contrary, computer simulations are almost always analogues of well-known animal research. A simulation should not be called an *experiment*, because it produces no new data; rather, it produces predictions of the data that should be observed in animals if the theory being tested is correct. The common tendency to call them experiments probably arises because they seem to have a subject and an environment, and the outcome is unknown to the researcher at the beginning. The parallel seems even stronger when the model or training uses randomization, which is known to affect the outcomes significantly in many cases. But everything on the computer is deterministic: Even "random" events are actually pseudorandom. The outcomes of simulations are completely determined and perfectly replicable. They cannot substitute for experiments on the behavior of intact organisms.

*Do Behavioral Models Need Neural Plausibility?*

I believe that Donahoe et al.'s approach to computer modeling with an emphasis on neurological plausibility is valuable, and in the long run essential. However, this does not imply that it is the only valid or even preferred way other behavior analysts should approach modeling. On the contrary, it is premature to impose a criterion of neurological plausibility; doing so forces working with incomplete models of behaving systems, which in turn can lead to erroneous conclusions. These points are readily illustrated by two examples from Donahoe et al.

---

Correspondence concerning this article should be addressed to William R. Hutchison, Behavior Systems, LLC, 2002 Linden Drive, Boulder, Colorado 80304 (E-mail: whutchi@behaviorsys.com).

The authors overstate our understanding of the brain. Even as I was writing these comments, I listened to a distinguished neuroscientist express "bafflement" that anyone would depend on current knowledge of the brain for behavioral models. Even individual neurons and synaptic processes are extraordinarily complex, apart from the complexity of the brain as a system, and in a complex system, even revising the view of how a single element works can drastically change the view of how the rest of the system works.

The authors have previously made a pragmatic argument for imposing neural plausibility (Donahoe, Burgos, & Palmer, 1993), pointing out how Darwin's theory of evolution was not widely accepted until the mechanisms of genetics were discovered. They argued that biological plausibility of our learning models was analogously important to persuade critics that there is a mechanism underlying it. However, the analogy breaks down because there is widespread acceptance of the brain as a sufficient mechanism for learning, even if we do not understand the details very well. On the other hand, there is great skepticism that learning can account for complex behavior such as language, an issue for which behavioral modeling is better suited to resolve.

Computer modeling of behavior can be done using purely behavioral formulations such as that of Rescorla and Wagner (1972) and behavioral economics (see *JEAB*, Vol. 64, pp. 257–431, 1995) rather than using neurologically based assumptions. Donahoe and Palmer (1994, p. 149) acknowledge that "The considerations that should inform computer simulations is a matter of current controversy," but Donahoe et al. conclude that "to be congenial with behavior analysis, all intraorganismic events must be the product of *independent* biobehavioral research" (p. 203). That would be true if behavior-analytic modelers wanted to claim that the implementation details of their models had the status of intraorganismic events, but that claim is generally disavowed and of secondary interest at most. Behavioral theorists and modelers already have a daunting enough challenge to produce comprehensive behavioral models, and find behavior-level modeling very congenial. Donahoe et al. cite Skinner's point (1938, p. 432) that nothing learned about the

physiology of behavior can ever undermine valid behavioral laws. We can turn the tables even further and assert that any model that is not behaviorally plausible cannot be neurally plausible. We have learned a great deal about adaptive system processes from work with neural network models that were developed without concern for biology or behavior, including the most common model in the field, backpropagation, which was invented as a mathematical method to extend linear regression (Werbos, 1974).

Because they restrict themselves to neurally plausible elements, Donahoe et al.'s model (their Figure 2) is an incomplete subsystem. One obvious missing element is a function relating the dopamine injections to occurrences of primary and secondary consequences of various kinds, magnitudes, and delays, that is, behavioral economic relations. Because the neural data they cite require that the injection occur within 200 ms, there is a significant challenge to reconcile that with the timing of consequences in the real world. Other missing elements seem to be necessary, such as a mechanism for the system's responses to function as stimuli even when the responses do not change the environment in a perceptible way.

I raise these examples of the incompleteness of Donahoe et al.'s model not to criticize what they have done, but to focus attention on an important consequence of a strict neural plausibility criterion. Subsystems are valuable, but some of the conclusions Donahoe et al. reach cannot be made on the basis of a subsystem. Specifically, although two of their conclusions are correct for the subsystem, they are invalid for complete behavioral system models and, by implication, for animals. These two elements of the simulation are especially relevant because behavior analysts who use computer models have to make practical choices about both of them.

#### *Are Linear Models Wrong?*

The authors argue (p. 203) that direct sensory-to-motor associations are not observed in brains. However, direct control mechanisms evolved very early and have many advantages: Learning is far simpler, much faster, requires far fewer resources, and usually shows more robust generalization to new situations. Animals learn more quickly on most problems

than hidden-layer neural network models do, with acquisition processes that in my judgment resemble those of linear models much more than hidden-layer models. Another line of evidence is the success of linear modeling in real-world relationships, which has been extremely successful (Dawes, 1979), with little need in most cases for the addition of nonlinear elements. It seems implausible that an evolutionary process would lead to the loss of such a simple and effective mechanism and the imposition of intermediate layers (two in the Donahoe et al. model) for *all* learning, with these huge disadvantages. Far more likely would be the addition of mechanisms for combinatorial learning into the existing, very effective architecture.

We could take our skepticism even further and question whether hidden units are necessary at all in an adaptive system. The authors acknowledge that linear networks can learn a surprising range of relationships, but they agree with most theorists that Minsky and Papert (1969) proved that hidden units were necessary for certain kinds of situations, including the prototypical combinatorial problem, exclusive-OR (XOR) (p. 203). However, Minsky and Papert proved their point only for a subsystem limited in critical ways like Donahoe et al.'s, whereas I have long argued that there are behavior-level theories that are sufficient to explain how an operant system could solve such problems without hidden nodes. I simulated two of these formulations specifically for the present commentary.

The first simulation (Hutchison, 1997, XOR) shows how response-produced stimuli enable a linear network to behave in multiple time steps like a network with hidden nodes. I added four additional responses to Minsky and Papert's "yes" and "no" responses, and gave the system sensors activated by the emission of its own responses. Any four responses would work, but the simulation used motor responses, which would tend to occur naturally upon presentation of spatially separated stimuli. I presented the four XOR cases in random order, and the system learned to emit one of the motor responses to each stimulus pattern (e.g., move left when only the left stimulus was presented, move forward when both stimuli were presented) as its initial response. The next time step, the system

emitted the correct yes-no response with perfect accuracy under control of the stimulus produced by its prior response. The function of hidden layers was performed by repeated activity of the linear system.

The second simulation (Hutchison, 1997, Match) showed yet another behavioral process that solves not only XOR but a much broader class of nonlinear problems such as matching to sample (MTS) and delayed MTS within a linear network. The central behavioral formulation of this simulation is closely related to Lowenkron's concept of *joint control* (1991), which is a special case of combining separately conditioned antecedents (Rescorla & Wagner, 1972) if we accept the equivalence of respondent and operant conditioning, as Donahoe et al. advocate. Unlike Donahoe et al.'s passive subsystem, a behaving system does not have to sense both (more generally, all) the stimuli simultaneously and respond instantaneously. It can, and often must, sense the stimuli in sequence (e.g., looking at one after another) and respond to them in sequence, sometimes with a delay. Its own responses become part of the stimuli that control subsequent responses. Lowenkron discusses verbal tact and self-echoic responses, but the analysis applies more generally to any discriminated and self-repetitious responses and to many types of stimuli and experimental arrangements.

I trained a simulated organism with no hidden units to tact each of four different stimuli, and then to self-echo the tact if a new stimulus was not presented. The eight combinations of two pairs of the stimuli (AA, AB, BA, BB, CC, CD, DC) were then presented in random order. In each case, the system tacted the first stimulus presented (e.g., A), after which the second stimulus was presented (A or B). If the second stimulus was the same as the first, the same response was emitted with increased strength due to the joint stimulus control as a tact and self-echoic. Conversely, when the two stimuli were different, the new external stimulus was an  $S^A$  for the self-echoic response (i.e., tacting A was weakened by presentation of the stimulus for B due to prior consequences for that relation), whereas the stimulus produced by the system's own previous tact was an  $S^A$  for the second tact. This reliable difference in strength transitions always occurred before the system was required

to emit its response indicating same or different, so reinforcement of correct same-different responses produced effective stimulus control by that event within 100 trials. When delays were introduced between presentation of the two stimuli, the system still produced the correct answers by emitting self-echoes during the delay, as we assume humans often do.

This process has important advantages over hidden-unit solutions. The system can almost immediately discriminate matches for any new pairs of stimuli it can tact. No additional training is necessary, because the stimulus control is by a property that is independent of the particular stimuli. This process seems to help explain many instances of complex behavior that are otherwise difficult to explain, such as *parity* (Palmer, 1996).

These two demonstrations do not prove that combinatorial nodes and neurons do not exist (they probably do), but the simulations thoroughly refute the claimed necessity of such elements in the prototypical case, XOR, and even in the more difficult delayed MTS case.

#### *Is Spontaneous Noise Necessary?*

The second point on which I believe working with a subsystem has led Donahoe et al. to an erroneous conclusion is that spontaneous activity is a *necessary* property of neurons (p. 208). At the beginning of my simulation work I programmed a mechanism for injecting spontaneous noise into my deterministic model, both because seemingly random variations exist in animal behavior and because theoretical analysis shows that noise is necessary to avoid behavior being trapped in locally optimal but globally suboptimal patterns. These local optimum problems do require noise, and for a subsystem such as Donahoe et al.'s it must be added externally, as they concluded. However, in a complete behavioral model the variability in real-world data combines with the variability produced by the many dynamic (chaotic?) processes of the environment-behavior system to produce variability that resembles spontaneous noise. The dynamic effects include response-produced environmental changes, orienting response-produced sensory changes, antecedent control by previous responses, and continuous changes in behavioral relations

from learning—all mentioned by Donahoe et al. All my simulations (Hutchison, 1997) show this variability, which anyone replicating them can track to its specific sources by examining the behavioral relations inside the system—a luxury not yet available with living subjects. The system's behavior has exhibited enough of this variability to solve a wide range of problems without adding any random noise to its deterministic behavior. Of course, living organisms *do* get trapped by local optima in some conditions, so our challenge is not to produce optimal models but models that behave like animals. Some intracellular noise may occur in living systems, but the necessity of spontaneous noise in behavioral process has not been proven and will probably be difficult to prove.

#### *Problems with Emphasizing Neural Plausibility*

From these examples, there are two drawbacks of emphasizing biological plausibility. First, it restricts practitioners to relatively simple subsystems, from which wrong conclusions may be drawn. Second, it may provide misleading suggestions about functional relationships. For example, I do not doubt Donahoe et al.'s assertion that in human brains, sensory neurons do not connect directly to motor neurons. But the apparent model design recommendations from this anatomical datum may be incorrect. This does not diminish the value of Donahoe et al.'s emphasis; rather, it affirms it. Only researchers with their joint behavioral and biological emphasis would be likely to notice the inconsistency, and it would direct their research very specifically. In the example of direct connections, they might look for ways in which the brain manages to behave as though it had direct connections. This would happen if many of the sensory association neurons were activated by simple sensory inputs and many of their motor association neurons activated simple responses, thereby mimicking a linear system.

I encourage theorists to use computer models to ask any kind of "what if?" question that seems to be worth the effort. Some may be concerned that this freedom from neural constraints, like freedom of speech, might degrade into anarchy. However, as with free speech, the audiences for published simula-

tions can judge for themselves what conclusions should be drawn from them. In our current state of knowledge, it is usually more productive to suspend judgment about the plausibility of assumptions and instead focus on the plausibility of the conclusions they produce in simulations—usually a much easier assessment. If the conclusions (the behavior in various situations) are implausible, then the assumptions as a set are implausible.

Behavior analysis does not have a complete system model that incorporates all the environment-behavioral relations we have discovered to predict accurately the range of behavior we observe. We can hope that neuroscience will provide suggestive analogies and a basis for preferring certain elements of our model, but we cannot afford to be restricted by dubious neural plausibility considerations at this stage. Only after we have more than one effective behavioral model will we have the need—or the luxury—to choose among them. Such models will be of great value to both behavior analysts and neuroscientists.

#### REFERENCES

- Dawes, R. M. (1979). The robust beauty of improper linear models in decision making. *American Psychologist*, 34, 571–582.
- Donahoe, J. W., Burgos, J. E., & Palmer, D. C. (1993). A selectionist approach to reinforcement. *Journal of the Experimental Analysis of Behavior*, 60, 17–40.
- Donahoe, J. W., & Palmer, D. C. (1994). *Learning and complex behavior*. Boston: Allyn & Bacon.
- Hutchison, W. R. (1984, December). Cognitive vs. behavioral psychology in artificial intelligence. *Newsletter for Behavioral Artificial Intelligence*, pp. 5–20.
- Hutchison, W. R. (1985, March). Behavior analysis and other adaptive systems theories within the “functional camp” of artificial intelligence. *Newsletter for Behavioral Artificial Intelligence*, pp. 3–18.
- Hutchison, W. R. (1995, December). *Teaching an agent to speak and listen with understanding: Why and how?* Proceedings of the Intelligent Information Agents Workshop, CIKM, Baltimore. (Permanent online archive at <http://www.cs.umbc.edu/~cikm/iaa/submitted/viewing/whutchi.html>)
- Hutchison, W. R. (1997). *Simulations of operant behavior: Software and documentation*. Cambridge, MA: Cambridge Center for Behavioral Studies.
- Lowenkron, B. (1991). Joint control and the generalization of selection-based verbal behavior. *The Analysis of Verbal Behavior*, 9, 121–126.
- Minsky, M. L., & Papert, S. A. (1969). *Perceptrons*. Cambridge, MA: MIT Press.
- Palmer, D. C. (1996). Achieving parity: The role of automatic reinforcement. *Journal of the Experimental Analysis of Behavior*, 65, 289–290.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning* (Vol. 2, pp. 64–99). New York: Appleton-Century-Crofts.
- Skinner, B. F. (1938). *The behavior of organisms*. Englewood Cliffs, NJ: Prentice Hall.
- Stephens, K. R., & Hutchison, W. R. (1993, October). Behavior analysis and the quest for machine intelligence. *Educational Technology*, pp. 52–61.
- Werbos, P. J. (1974). *Beyond regression: New tools for prediction and analysis in the behavioral sciences*. Unpublished doctoral dissertation, Harvard University.